

Claims:

1. A method for providing a data management system, comprising:
preprocessing a database having a relation to produce an index;
receiving a query having aggregation constraints; and
applying said index to look up a result in response to said query having aggregation constraints.
2. The method of claim 1, wherein said query having aggregation constraints includes an n-dimensional vector of constants associated with said aggregation constraints.
3. The method of claim 1, wherein said index is produced by evaluating solution to a knapsack problem on said relation for a select number of vectors of constants associated with said aggregation constraints.
4. The method of claim 1, wherein said index contains pointers pointing to one or more answers that are considered to be said result within a predefined approximation factor.
5. The method of claim 1, wherein said index contains pointers pointing to a plurality of partitions, where if said query falls within one of said partition, then each partition is representative of a set of answers corresponding to said result.
6. The method of claim 1, wherein said preprocessing step comprises:
identifying a dominating vector of constants, \bar{c}' for a given n-dimensional vector of constants \bar{c} .
7. The method of claim 6, wherein said preprocessing step further comprises:
obtaining a partition defined by said vector \bar{c} and said vector \bar{c}' .

8. The method of claim 7, wherein said partition is expressed as a hyper rectangle.

9. The method of claim 7, wherein said preprocessing step further comprises:

inserting said partition into a multidimensional data structure.

10. The method of claim 9, wherein said multidimensional data structure is an R-Tree.

11. The method of claim 1, wherein said result is guaranteed to be accurate within a predefined approximation factor.

12. The method of claim 11, wherein said predefined approximation factor can be selectively changed.

13. The method of claim 1, wherein said result is representative of one of more answers that are deemed to be dominant.

14. The method of claim 1, wherein said result is representative of one of more answers that are deemed to be dominant within an approximation factor.

15. A method for generating an index for use with query having aggregation constraints, comprising:

identifying a dominating vector of constants, \bar{c}' for a given n-dimensional vector of constants \bar{c} ;

obtaining a first partition defined by said vector \bar{c} and said vector \bar{c}' ; and inserting said first partition into a multidimensional data structure.

16. The method of claim 15, further comprising excluding said first partition; and

wherein said steps are iteratively repeated until an entire feasible space dominated by said vector of constants \bar{c} has been partitioned.

17. An apparatus for providing a data management system, comprising:
means for preprocessing a database having a relation to produce an index;

means for receiving a query having aggregation constraints; and
means for applying said index to look up a result in response to said query having aggregation constraints.

18. An apparatus for generating an index for use with query having aggregation constraints, comprising:

means for identifying a dominating vector of constants, \bar{c}' for a given n-dimensional vector of constants \bar{c} ;

means for obtaining a first partition defined by said vector \bar{c} and said vector \bar{c}' ; and

means for inserting said first partition into a multidimensional data structure.

19. A computer-readable medium having stored thereon a plurality of instructions, the plurality of instructions including instructions which, when executed by a processor, cause the processor to perform the steps comprising of:

preprocessing a database having a relation to produce an index;
receiving a query having aggregation constraints; and
applying said index to look up a result in response to said query having aggregation constraints.

20. A computer-readable medium having stored thereon a plurality of instructions, the plurality of instructions including instructions which, when executed by a processor, cause the processor to perform the steps comprising of:

identifying a dominating vector of constants, \bar{c}' for a given n-dimensional vector of constants \bar{c} ;

obtaining a first partition defined by said vector \bar{c} and said vector \bar{c}' ; and inserting said first partition into a multidimensional data structure.